INTRODUCTION TO CLOUD SYSTEMS

Lecture 9 – Security Development Lifecycle, Confidential computing: SGX, TME usage in Cloud, Out of band management – configuration, telemetry

Security Development Lifecycle

Developing products with a security mindset is an important industry practice that reduces mitigation costs and improves product resiliency. The security development lifecycle (SDL) is a set of processes that implement security principles and privacy tenets into product development. These processes incorporate security minded engineering and testing at the onset of product development when it is more effective and efficient to employ.

Security Development Lifecycle

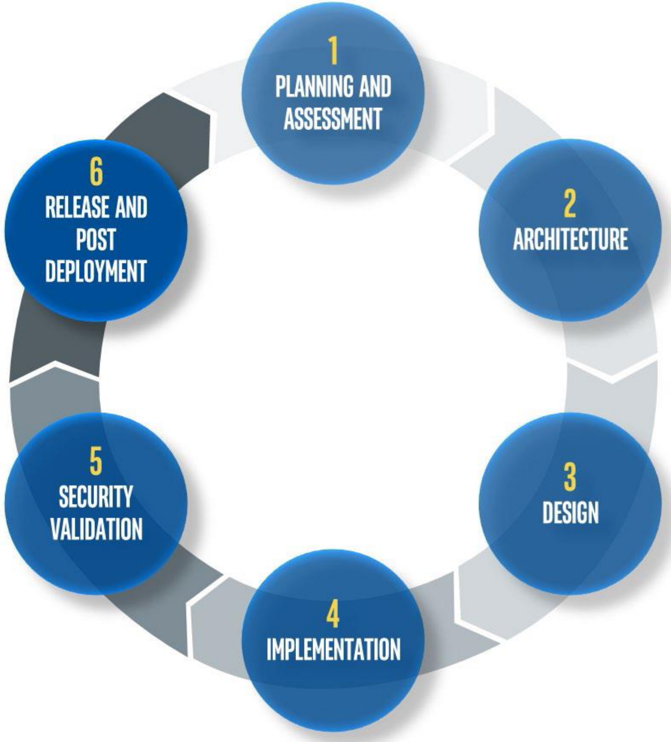

Planning and assessment: This first step in the SDL process helps the project team identify what tasks and activities will be needed throughout a project’s development lifecycle. These tasks are specifically assigned to individual projects and are tailored to the project’s expected security and privacy risks by security experts assisting product architects.

https://newsroom.intel.com/wp-content/uploads/sites/11/2020/10/sdl-2020-whitepaper.pdf

Security Development Lifecycle

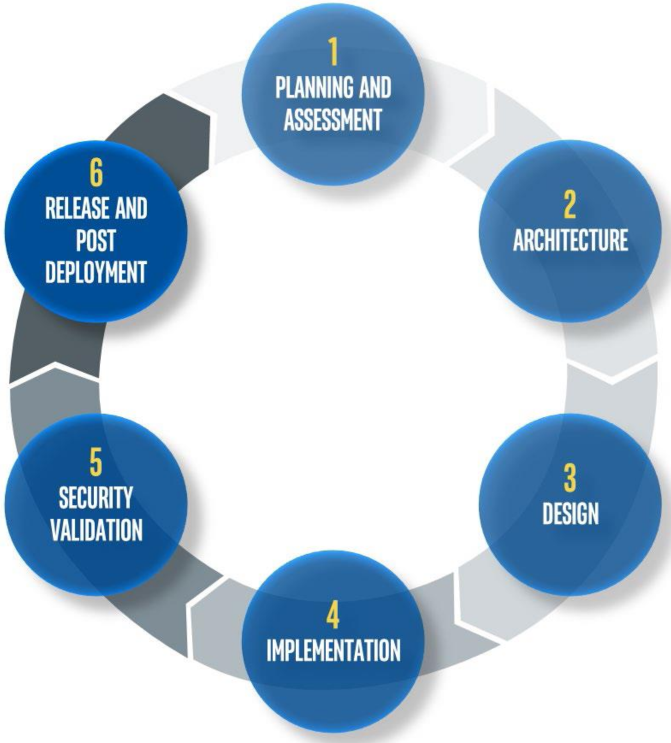

Architecture: Architects and developers collaborate to define security objectives and use them to build an appropriate threat model. Following industry best practices for secure design principles, the team executes the reviews laid out in the planning and assessment phase, documenting all work. If appropriate, the team will complete a series of architectural reviews during this timeframe, including cryptography and more.

https://newsroom.intel.com/wp-content/uploads/sites/11/2020/10/sdl-2020-whitepaper.pdf

Security Development Lifecycle

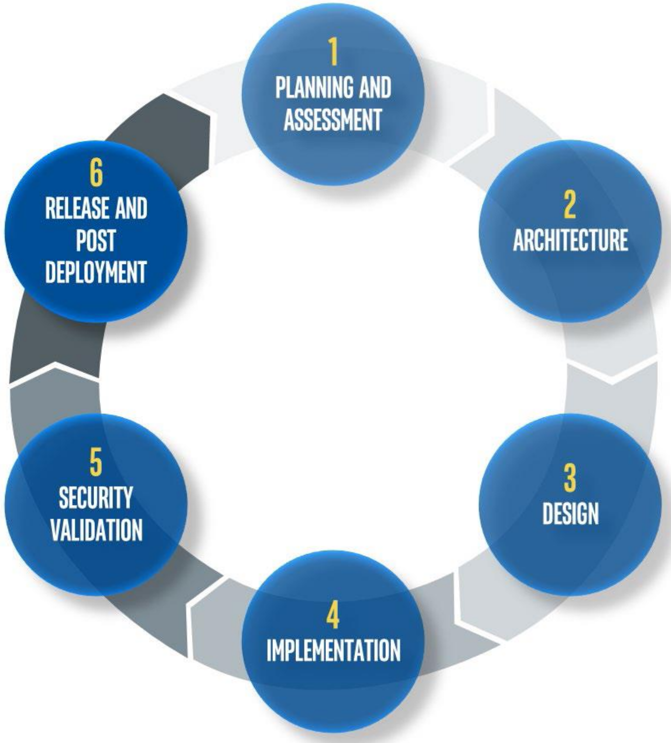

Design: Engineers perform design level security and privacy analysis based on the security objectives, threats and requirements identified in the previous phase. The team translates these items and updates security documentation. They define the security and privacy validation strategy in a manner intended to ensure that sufficient resources will be available to cover all requirements during the validation phase.

https://newsroom.intel.com/wp-content/uploads/sites/11/2020/10/sdl-2020-whitepaper.pdf

Security Development Lifecycle

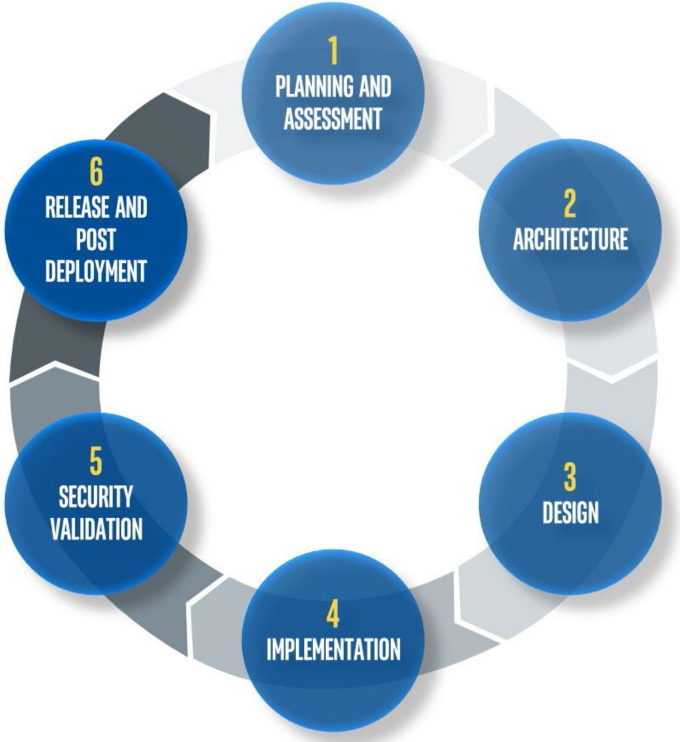

Implementation: The team works to ensure that the product implementation addresses the threat models defined in the architecture phase. Engineers perform secure code reviews and static code analysis, establish formal verification for any mitigations applied for potential vulnerabilities, and check that the architecture and design of the product is performing as intended. The team accounts for any necessary updates to the SDL and formal security validation plan to execute in the next phase.

https://newsroom.intel.com/wp-content/uploads/sites/11/2020/10/sdl-2020-whitepaper.pdf

Security Development Lifecycle

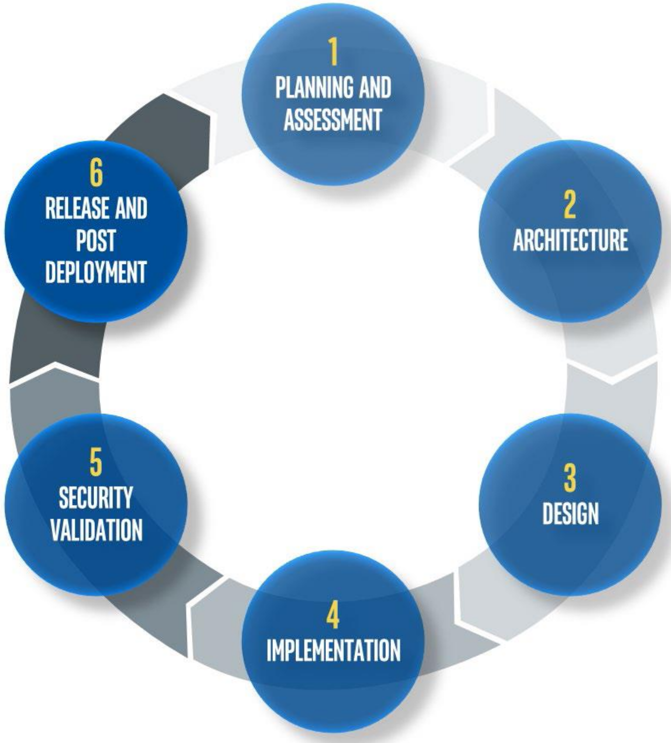

Security tools are essential to identifying potential and known vulnerabilities, and doing so at scale. That’s why Intel maintains a dedicated team for assessing and implementing security tools across the enterprise. There are also engineers developing custom fuzzing tools, among others, across the company for internal, academic, and customer use.

https://newsroom.intel.com/wp-content/uploads/sites/11/2020/10/sdl-2020-whitepaper.pdf

Security Development Lifecycle

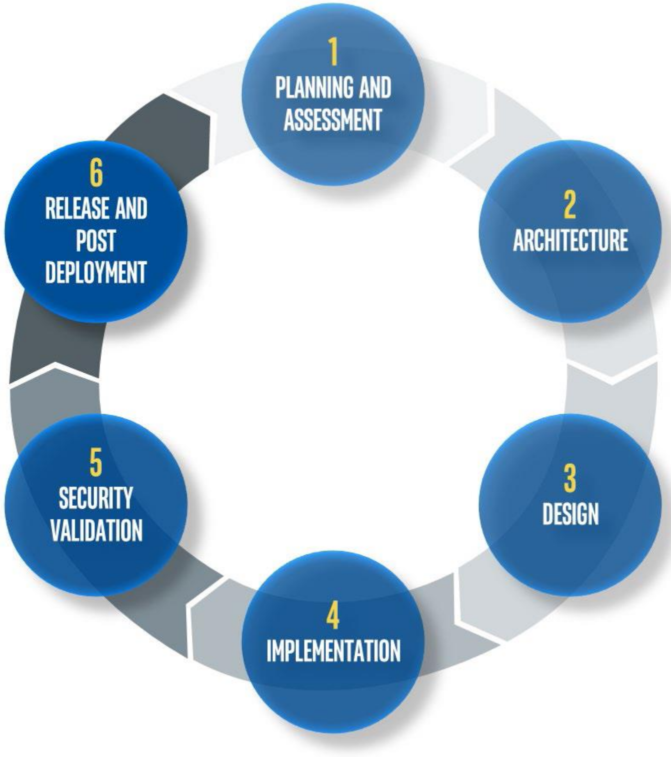

Security validation: When done well, validation monitors known vulnerabilities and emerging threats. It can even help anticipate potential threats. Ultimately, validation results affect “ship/no ship” decisions based on whether the product requirements have been met and required plans have been executed. Validation teams use many different forms of testing, including but not limited to, penetration testing (external if required) and fuzzing.

https://newsroom.intel.com/wp-content/uploads/sites/11/2020/10/sdl-2020-whitepaper.pdf

Security Development Lifecycle

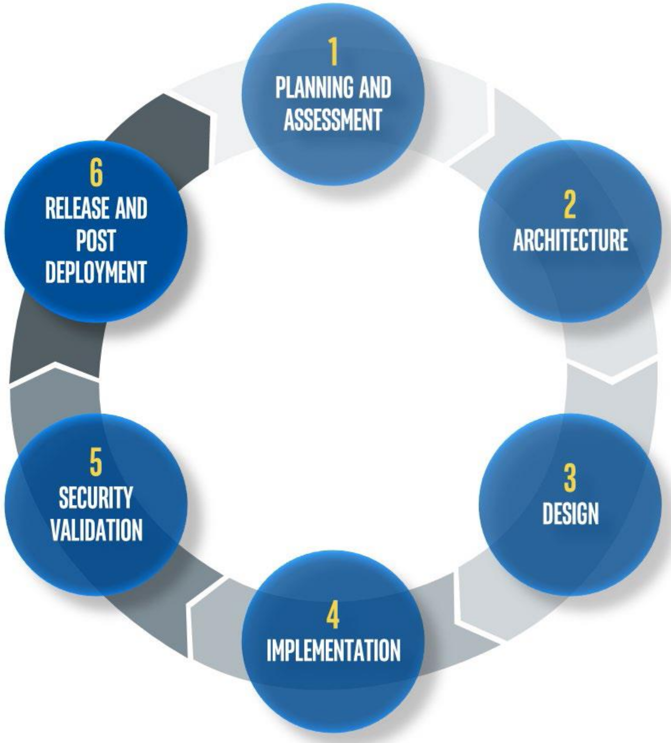

Release and post deployment: During this final phase, teams execute another round of testing and verify that previous issues were resolved. Teams scan for malware and for known vulnerabilities in third-party components and IP, updating them if appropriate. Finally, teams put a plan in place for product support and survivability. This includes vulnerability management over the lifespan of the product and triage/mitigation in partnership with Intel’s Product Security Incident Response Team (PSIRT).

https://newsroom.intel.com/wp-content/uploads/sites/11/2020/10/sdl-2020-whitepaper.pdf

Confidential computing

Almost all businesses within all sectors, and an increasing number of public and governmental institutions, move their data storage and processing onto cloud computing platforms.

- Better services,
- Scalability cost
- Efficiency
- Accessibility

Confidential computing

The **security is easier to maintain** and is significantly improved in well-designed cloud environments, where security protocols and encryption are built into the entire infrastructure by default, compared to most on-premises environments.

Confidential computing

User authentication and encryption are two of the key tools developers use for securing data in the cloud.

Data exists in three states: **In transit, at rest (stored) and in use.**

However, encryption has only been technically possible on a large scale for data in transit and when stored, and not when the data are used (during data processing in computer RAM). Encryption makes the data useless and is fundamental to preserve data privacy and security.

Confidential computing

To strengthen the security and “compensate” for the lack of encryption of the data throughout the entire data flow, contractual and organisational measures are a part of the equation. For processing of sensitive data, this may represent an unacceptable reliance on non-technical measures. As long as there is a theoretical possibility that the cloud provider may access and read your or my data, there will be some level of uncertainty.

Confidential computing

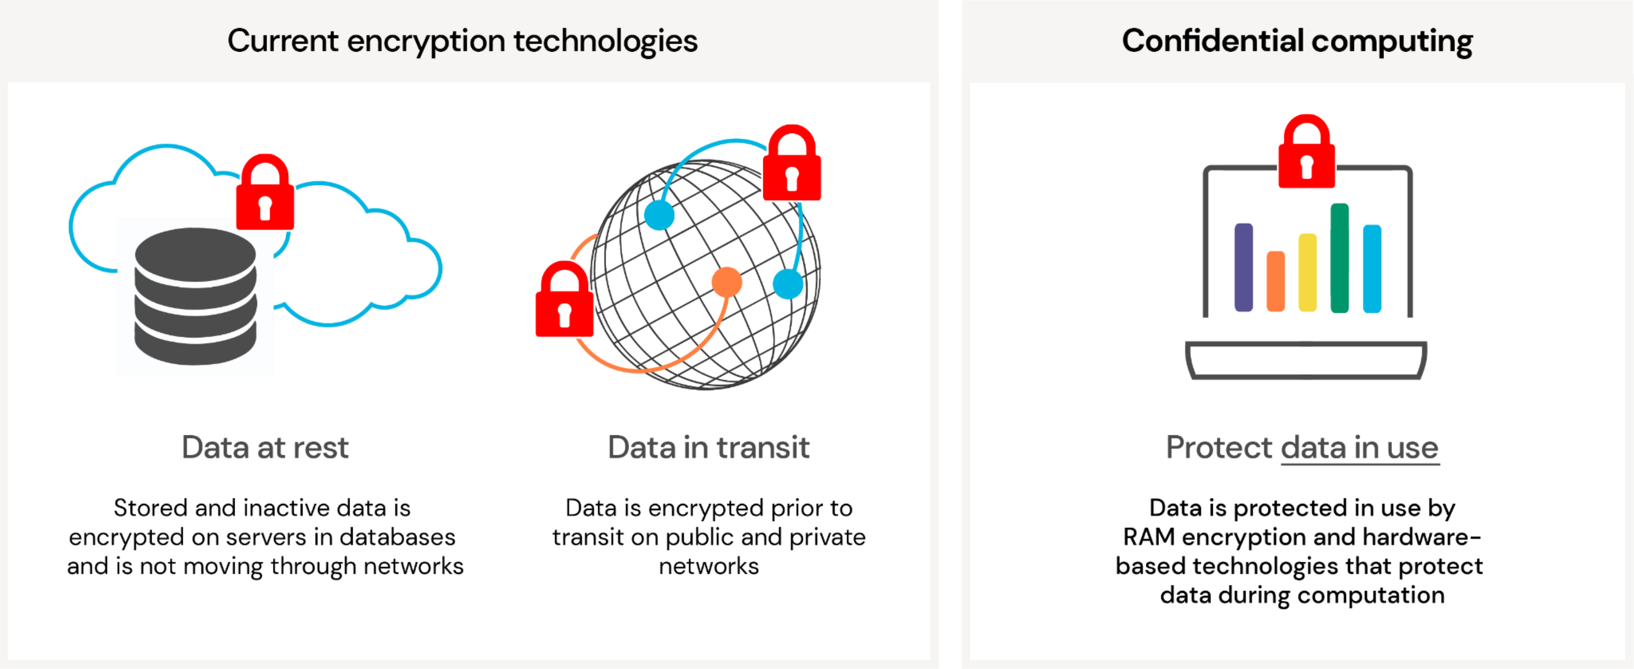

Current encryption technologies

Data at rest

Stored and inactive data is encrypted on servers in databases and is not moving through networks

Data in transit

Data is encrypted prior to transit on public and private networks

Confidential computing

Protect data in use

Data is protected in use by RAM encryption and hardware-based technologies that protect data during computation

Confidential computing

Protecting the data in use is critical to offer complete security across the data lifecycle – Confidential Computing

Technologies that “isolate” sensitive data during processing from the cloud provider and all unauthorized personnel. The aim is to significantly improve the security and privacy of cloud computing and establish a hardware-based root of trust of the data processing. By making the data in use inaccessible or completely uninterpretable by encryption, the reliance on organisational and contractual commitments is less important.

Confidential computing

There have been several definitions of the term “Confidential computing” from when it was introduced, but regardless of the wording, the essence is that:

“Confidential Computing protects data in use by performing computation in a hardware-based Trusted Execution Environment (TEE). These secure and isolated environments prevent unauthorized access or modification of applications and data while they are in use, thereby increasing the security level of organisations that manage sensitive and regulated data”.

Confidential computing

The TEE, also called an enclave, executes the code and isolates and protects the code and data from the host system (plus the host system’s owners, e.g. the cloud provider), and may also provide code integrity and attestation that can be used for compliance and regulatory purposes. **In the TEE, the data is secure in the memory, when traveling to and from the host CPU, and finally, during execution on the host CPU.**

Confidential computing

Now „Confidential computing” is available for all to use. This means that the entire flow of data in a well-designed cloud environment is protected; from login with authentication mechanisms such as two-factor authentication, to encryption of data in transit, encryption of stored data, and now finally and critically, protection of data in use by memory encryption or other hardware-based technologies

Confidential computing

Hardware-based root of trust

A hardware-based root of trust is the basis for confidential cloud computing. In a trust hierarchy, the hardware itself is the foundation that all the other layers of trust are laid upon, such as the BIOS (Basic Input/Output System), operating system, firewalls, security software, physical security, and ultimately organisational and contractual measures.

Confidential computing

Hardware-based root of trust

If the data in use is protected by the hardware itself, which is designed in a manner that is inherently resistant to malware injection and tampering, the memory is encrypted, and the hardware cannot be accessed in any way by any ports or APIs by any unauthorized persons including the cloud operators, then we are able to reach our security objective which is to protect the data end-to-end: In transit, at rest and in use. This is now becoming the reality.

Confidential computing

Intel® SGX

Intel® SGX, by bypassing a system’s operating system (OS) and virtual machine (VM) software layers, provides significant additional protection against many of these kinds of attacks and adds data security and addresses the need for more confidential computing. It provides a hardware-based security solution that utilizes encryption to change how memory is accessed, providing enclaves of protected memory to run your application and its data.

Confidential computing

Intel® SGX

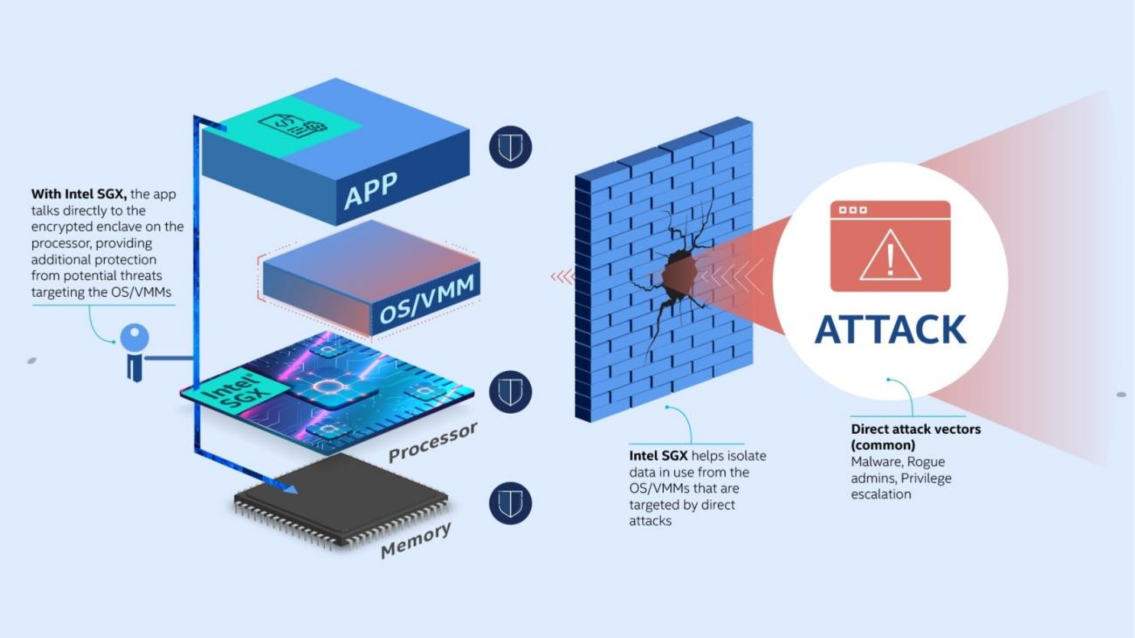

With Intel SGX, the app talks directly to the encrypted enclave on the processor, providing additional protection from potential threats targeting the OS/VMMs

Intel SGX helps isolate data in use from the OS/VMMs that are targeted by direct attacks

Direct attack vectors (common) Malware, Rogue admins, Privilege escalation

Confidential computing

Intel® SGX

Side-channel attacks

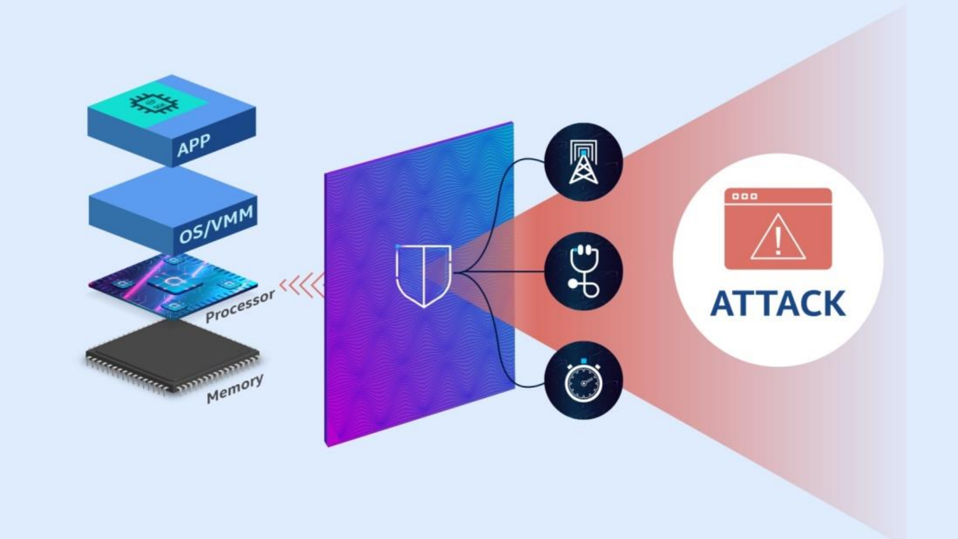

Side-channel attacks are based on using information such as power states, emissions and wait times directly from the processor to indirectly infer data use patterns. These attacks are very complex and difficult to execute, potentially requiring breaches of a company’s data center at multiple levels: physical, network and system.

Confidential computing

Intel® SGX

Intel Software Guard Extensions (Intel SGX) provides fine grain data and privacy protection via application isolation in memory, independent of operating system or hardware configuration.

Confidential computing

Intel® TME

Intel TME (Total Memory Encryption) encrypts all data passing to and from a computer's CPU with a single transient key. Such information includes customer credentials, encryption keys, and other IP or personal information.

Memory attacks have quietly emerged as a new class of hacking techniques to undermine conventional security measures.

This new threat includes attacks at the hardware level such as removal and reading of dual in-line memory modules (DIMMs) or the installation of attack hardware.

Without Intel TME, hackers can access critical data, encryption keys, or install malware, compromising the security of a system.

Confidential computing

Intel® TME – how it works?

Intel TME begins in the early stages of the boot process. Once configured and locked, it will encrypt all the data on the external memory buses of a CPU with the NIST Standard AES-XTS algorithm with 128-bit keys.

The encryption key is generated using a hardened random number generator in the CPU and never exposed to software, allowing existing software to run unmodified while better protecting memory. A new platform key is generated by the processor on every boot.

Confidential computing

Intel® TME – how it works?

Data in memory and on the external memory buses is encrypted and is only in plain text while inside the CPU, similar to storage encryption on hard disks or SSDs.

There are, however, some instances where it would be better to not encrypt a portion of memory, so Intel TME allows the BIOS to specify a physical address range to remain unencrypted. TME can be enabled or disabled by IT admins in the BIOS settings.

Out-of-Band Management

What Is Out-of-Band Management?

Out-of-band (OOB) management is a method of remotely controlling and managing critical IT assets and network equipment using a secure connection through a secondary interface that is physically separate from the primary network connection.

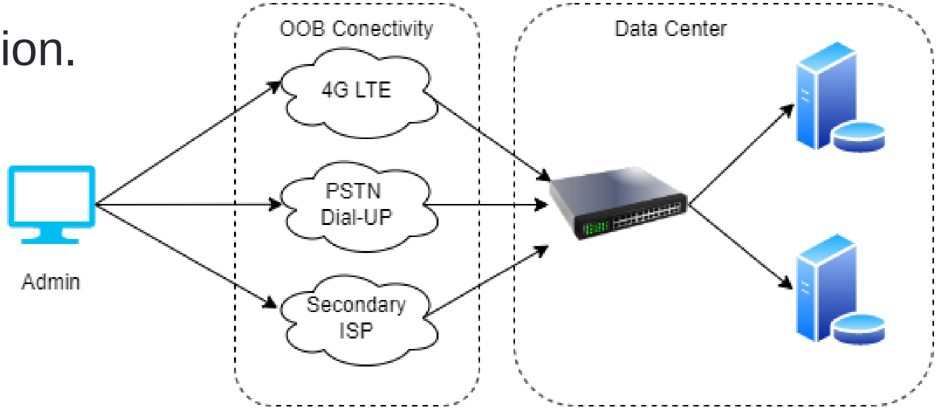

Out-of-Band Management

In-Band Management vs. Out-of-Band Management

In-Band Management (software based) allows IT administrators to manage devices from a single interface. Software-only solutions are limited, because devices must be powered on and capable of connecting to IT services before they can be managed.

With hardware-based out-of-band management (OOBM), IT administrators can access devices even if they are turned off or the operating system (OS) is down or unresponsive.

Out-of-Band Management

Benefits of Out-of-Band Management

Better access and functionality - out-of-band management capabilities are hardware based and they operate beneath the OS. Administrators can set levels on the BIOS and UEFI firmware interfaces or make elevated task changes. It is also possible to perform routine tasks, like PC setup and configuration and OS or security updates.

Out-of-Band Management

Benefits of Out-of-Band Management

Lower IT management costs - Businesses that have hundreds or thousands of PC-based devices in many locations can reduce their operational costs related to travel and IT staff hours.

Faster fixes - there’s no need to wait for machines to be shipped back to IT or for technicians to reach the site. Devices can be up and running faster, reducing downtime and boosting productivity.

Out-of-Band Management

IPMI

The Intelligent Platform Management Interface (IPMI) is a set of computer interface specifications for an autonomous computer subsystem that provides management and monitoring capabilities independently of the host system's CPU, firmware (BIOS or UEFI) and operating system. IPMI defines a set of interfaces used by system administrators for out-of-band management of computer systems and monitoring of their operation.

Out-of-Band Management

IPMI

The primary IPMI features include:

- Monitoring (supervision of the hardware)
- Recovery Control (Recover/Restart the server)
- Logging (protocol "out-of-range" states for the hardware)
- Inventory (list of hardware inventory)
- Available even if system is powered down and no OS loaded

Out-of-Band Management

BMC

The baseboard management controller (BMC) provides the intelligence in the IPMI architecture. It is a specialized microcontroller embedded on the motherboard of a computer – generally a server. The BMC manages the interface between system-management software and platform hardware. BMC has its own firmware and RAM.

Out-of-Band Management

BMC

Different types of sensors built into the computer system report to the BMC on parameters such as temperature, cooling fan speeds, power status, operating system (OS) status, etc. The BMC monitors the sensors and can send alerts to a system administrator via the network if any of the parameters do not stay within pre-set limits, indicating a potential failure of the system. The administrator can also remotely communicate with the BMC to take some corrective actions – such as resetting or power cycling the system to get a hung OS running again. These abilities reduce the total cost of ownership of a system.

Out-of-Band Management

BMC

Physical interfaces to the BMC include SMBuses, an RS-232 serial console, address and data lines and an IPMB, that enables the BMC to accept IPMI request messages from other management controllers in the system.

IPMI Block Diagram

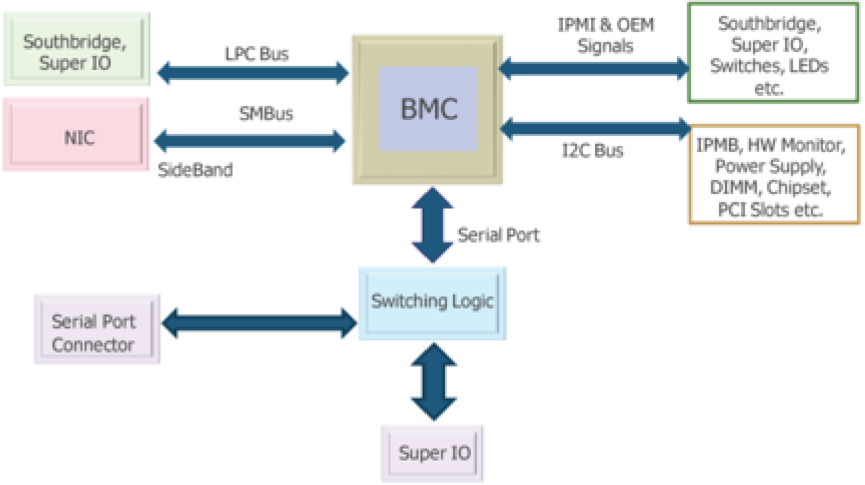

Out-of-Band Management

BMC

A direct serial connection to the BMC is not encrypted as the connection itself is secure. Connection to the BMC over LAN may or may not use encryption depending on the security concerns of the user.

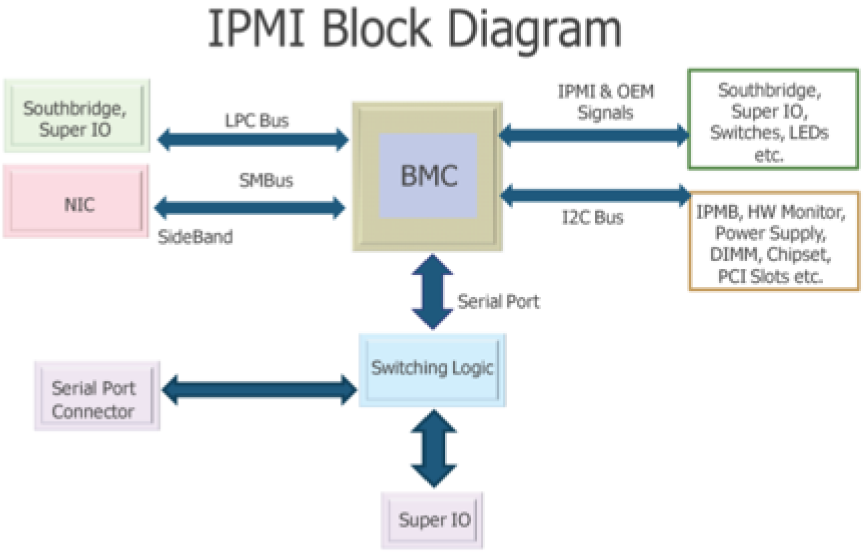

Out-of-Band Management

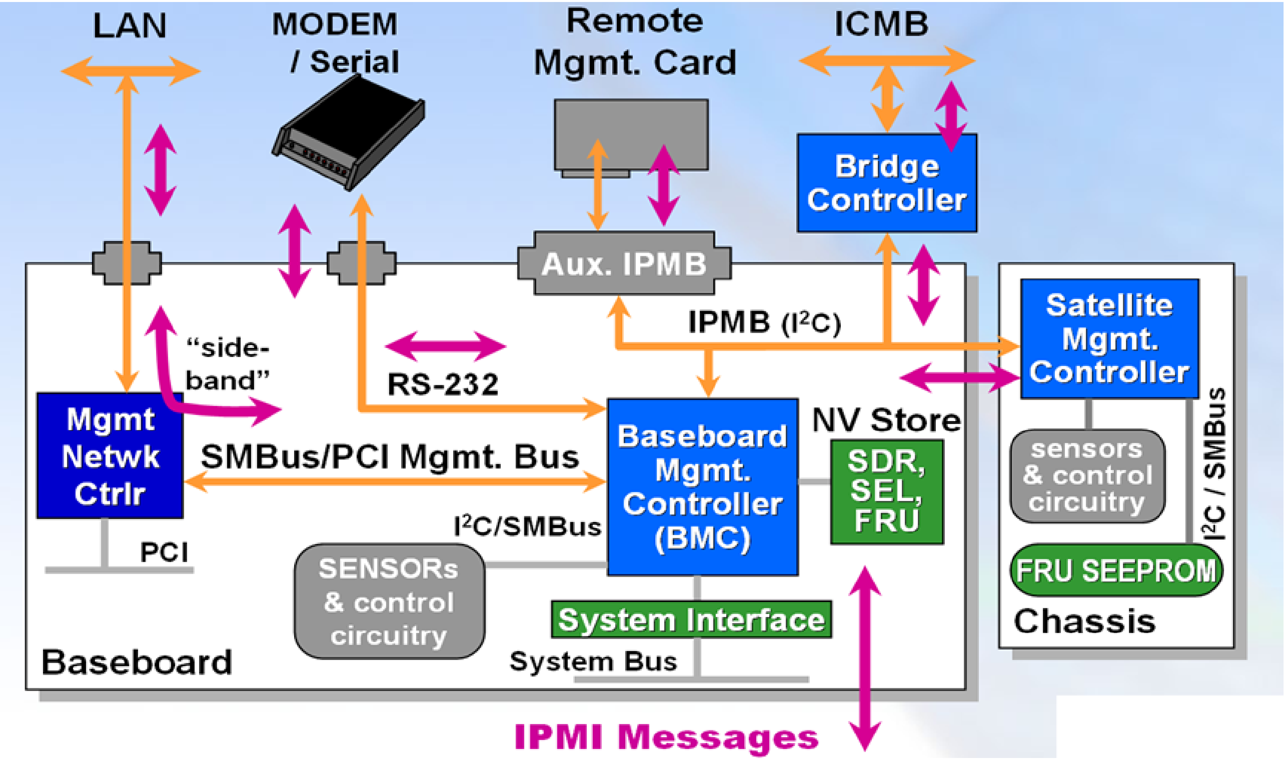

https://www.cern.ch/it-dep-fio-ds/Presentations/2004/ipmi_server_management.ppt

Out-of-Band Management

Redfish

- An open industry standard specifiction that provide simple, modern and secure management of scalable platform hardware.
- RESTful interface over HTTPs in JSON format bases on OData v4
- A secure, multi-node capable replacement for IPMI-over-LAN.

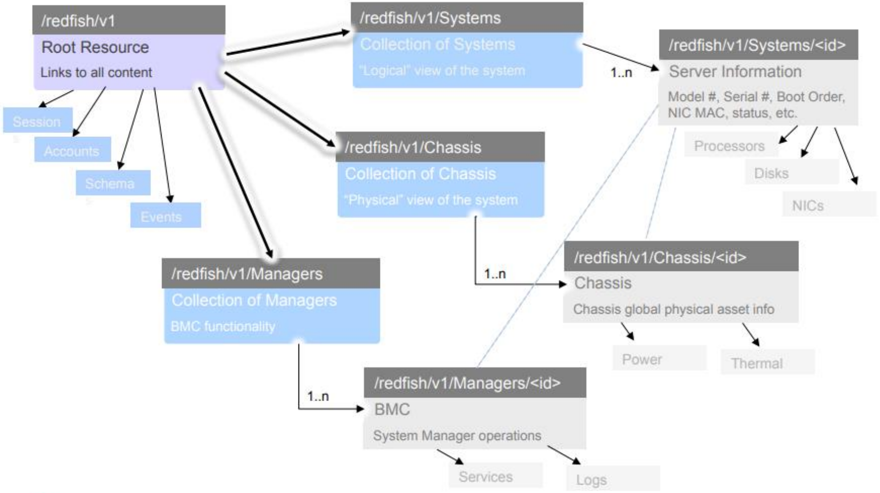

/redfish/v1
Root Resource
Links to all content

Session
Accounts
Schema
Events

/redfish/v1/Systems
Collection of Systems
"Logical" view of the system

/redfish/v1/Systems/<id>
Server Information
Model #, Serial #, Boot Order, NIC MAC, status, etc.

Processors
Disks
NICs

/redfish/v1/Chassis
Collection of Chassis
"Physical" view of the system

/redfish/v1/Chassis/<id>
Chassis
Chassis global physical asset info

Power
Thermal

/redfish/v1/Managers
Collection of Managers
BMC functionality

/redfish/v1/Managers/<id>
BMC
System Manager operations

Services
Logs

DMTF Redfish

www.dmtf.org